Matrix factoring by fraction-free reduction

Johannes Middeke
RISC Linz
Johannes Kepler University
Linz, Austria
jmiddeke@risc.uni-linz.ac.at

David J. Jeffrey
ORCCA and Dept Applied Mathematics
The University of Western Ontario
London, Ontario, Canada N6A 5B7
djeffrey@uwo.ca

arXiv:1603.03565v1 [cs.SC] 11 Mar 2016

ABSTRACT

We consider exact matrix decomposition by Gauss-Bareiss reduction. We investigate two aspects of the process: common row and column factors and the influence of pivoting strategies. We identify two types of common factors: systematic and statistical. Systematic factors depend on the process, while statistical factors depend on the specific data. We show that existing fraction-free QR (Gram–Schmidt) algorithms create a common factor in the last column of Q. We relate the existence of row factors in LU decomposition to factors appearing in the Smith normal form of the matrix. For statistical factors, we identify mechanisms and give estimates of the frequency. Our conclusions are tested by experimental data. For pivoting strategies, we compare the sizes of output factors obtained by different strategies. We also comment on timing differences.

Keywords

LU Decomposition; Fraction free; QR factors; Common factor removal; pivoting strategy

1. INTRODUCTION

Although known earlier, fraction-free methods for exact matrix computations became popular after Bareiss's study of Gaussian elimination [1]. Extensions to related topics, such as LU factoring, were considered in [9, 10, 15]. Gram–Schmidt orthogonalization and QR factoring were studied by [3], under the more descriptive name of exact division. Recent studies have looked at extending fraction-free LU factoring to non-invertible matrices [7] and rank profiling [2], and more generally to areas such as the Euclidean algorithm, and the Berlekamp–Massey algorithm [8]. We consider matrices over an integral domain $\mathbb{D}$. For the purposes of giving illustrative examples and conducting computational experiments, matrices over $\mathbb{Z}$ and $\mathbb{Q}[x]$ are used, because the metrics associated with these domains are well established and familiar to readers. We emphasize, however, that the methods here apply for all integral domains, as opposed to methods that target specific domains, such as [5, 12].

The starting point for this paper is the fraction-free form for LU decomposition [7]: given a matrix A over an integral domain $\mathbb{D}$,

$$A = P_w L D^{-1} U P_c \ , \tag{1}$$

where L, D and U are over $\mathbb{D}$. L and U are lower and upper triangular and their diagonals contain the pivots of the Gaussian elimination; D is diagonal and contains products of the pivots. The permutation matrices P_w and P_c ensure that the decomposition is always a full-rank decomposition, even if A is rectangular or rank deficient. In addition to the usual indeterminacy due to varying pivot choices, the columns of L and the rows of U can be multiplied by common factors, which then appear also in D. We show in section 3 that this form can cover QR decomposition also.

Our first main result is for QR factoring. In this context, the orthonormal Q matrix used in floating point calculations is replaced by a Θ matrix, which is left-orthogonal, i.e. $\Theta^t\Theta$ is diagonal, but $\Theta\Theta^t$ is not. We show that for a square matrix A, the last column of Θ, as calculated by existing algorithms, is subject to an exact division by the determinant of A, with a significant reduction in size. This is an example of a systematic factor, being one inherent to the algorithm.

Systematic factors occur in several ways. The Bareiss algorithm uses exact division precisely to remove systematic factors; the Gram–Schmidt algorithm from [3] is another, where exact division removes systematic factors during the reduction. In addition to these, we add a different type of systematic factor: we show a relation between GCDs existing for the rows in matrices obtained from LU factoring, and entries in the Smith normal form of the same initial matrix.

We next consider statistical factors: ones which depend on the initial data. When LU and QR matrices are computed using current standard fraction-free algorithms, their rows and columns may contain common factors. We discuss their origins and show we can predict a significant proportion of them from simple considerations. Their presence influences aspects such as uniqueness. Specifically, for the basic decomposition (1), we show how common factors can be moved between the three matrices. We discuss when this is beneficial.

We next consider the role of pivoting in Gaussian reduction. Geddes *et al.* [4] comment "We also mention that when the entries of $A^{(k)}$ are not of uniform size, it may be worthwhile to interchange rows in order to obtain a smaller pivot at the next step". It is often said that whereas for *floating-point* Gaussian elimination the largest pivot should be cho-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISSAC'16 July 19–22, 2016, Waterloo, ON, Canada

© 2016 ACM. ISBN 123-4567-24-567/08/06…$15.00

DOI: 10.475/123_4

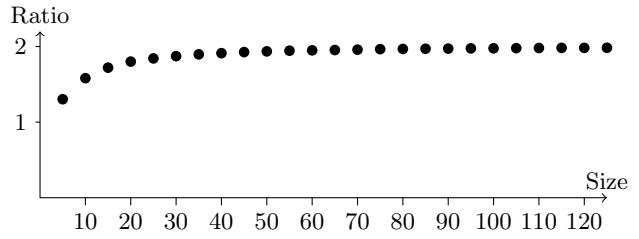

Figure 1: Comparison of the output size of Gaussian Elimination vs. $LD^{-1}U$. We show the ratio of the number of digits in the output of Gaussian elimination divided by the number of digits in the output of the $LD^{-1}U$ decomposition for random square integer matrices of various sizes.

sen, in the setting of exact computation the smallest pivot is best. Although within the floating-point literature, pivoting has been studied over an extended period, much less attention has been paid to the question in the context of exact computation. We consider a number of strategies empirically, and show that selecting the smallest pivot, suitably defined, leads to smaller output matrices.

The paper will start with a brief discussion of fraction-free methods, then present results for QR factoring, LU factoring, and finally pivoting.

2. FRACTION-FREE METHODS

Fraction-free methods are based on the assumption that it is more efficient to compute with the elements of the input domain of a matrix than to compute with elements from the quotient field. Since the solutions usually sought require the quotient field, fraction-free methods can be regarded as delaying for as long as possible the ultimate fall from grace of the computation. Here, some measurements are reported to supply empirical evidence to support fraction-free methods.

Our first point of comparison is between the LU decomposition offered by MAPLE, through `LUDecomposition(A)`, and our own implementation of the $LD^{-1}U$ decomposition based on [7]. The built-in MAPLE command returns matrices L and U such that all diagonal elements of L are 1, and both L and U contain elements from the quotient field of $\mathbb{D}$. The procedure which we implemented has the output format described in [7, Theorem 2].

Figure 1 shows the ratio of average storage requirements for the decomposition of integer matrices. Here, we measure the total number of digits needed to represent the final output. Note that this metric does not depend on the internal implementation of the two functions, nor does it depend on the particular computer algebra system. As figure 1 illustrates, fraction-free methods require roughly half the storage.

Table 1 compares timings for random integer matrices, as functions of matrix size and initial data size. For this experiment we used our own implementation of Gaussian elimination, since we do not know the details of MAPLE's built-in procedure, which may well use compiled code. By writing our own programme we make sure that every common part, for example pivot searching, uses exactly the same code and only the reduction steps differ. As table 1 reveals, the advantages of the fraction-free method are clear, while not spectacular.

s \ n	11	19	31	53	73	97	107
3	1.00	0.85	0.91	0.78	0.72	0.63	0.62
7	0.97	0.88	0.83	0.71	0.65	0.59	0.56
13	0.94	0.84	0.82	0.68	0.61	0.55	0.52
23	0.93	1.33	0.85	0.66	0.59	0.51	0.49
37	0.89	0.81	0.77	0.63	0.56	0.49	0.47
53	0.93	0.80	0.74	0.62	0.55	0.47	0.45
67	0.90	1.32	0.73	0.60	0.54	0.46	0.44
89	0.89	0.53	0.72	0.61	0.53	0.45	0.43
109	0.87	0.77	0.73	0.60	0.52	0.44	0.42

Table 1: Timings of $LD^{-1}U$ decomposition vs. Gaussian elimination for integer matrices. The table shows the run times for $LD^{-1}U$ divided by those for Gauss for random n-by-n matrices with maximal entry size s.

3. COMMON FACTORS IN QR

A fraction-free (exact division) algorithm for Gram–Schmidt orthogonalization was described by [3]. An algorithm based on LU factoring has been described in [13, 15]. The two approaches yield the same results. We denote the decomposition by $A = \Theta D^{-1} R$, because Q usually denotes an orthonormal matrix, and Θ is not orthonormal. We give a new statement of the basic theorem.

THEOREM 1. *Given a square, full-rank matrix A over an integral domain $\mathbb{D}$, the partitioned matrix $(A^t A, A^t)$ has a fraction-free LU decomposition*

$$(A^t A, A^t) = R^t D^{-1} (R, \Theta^t) \ ,$$

where $\Theta^t \Theta = D$ and $A = \Theta D^{-1} R$.

PROOF. We can apply LU factoring, to get

$$(A^t A, A^t) = \hat{L} D^{-1} (\hat{U}, \Theta) \ ,$$

where the notation $\hat{L}, \hat{U}$ emphasizes that the matrices refer not to a factoring of A, but of $A^t A$. Since this matrix is symmetric we obtain

$$\hat{L} D^{-1} \hat{U} = A^t A = \hat{U}^t D^{-1} \hat{L}^t .$$

Because A has full rank, so do L and U and we can rewrite the equation as

$$\hat{U} (\hat{L}^t)^{-1} D = D \hat{L}^{-1} \hat{U}^t .$$

Examination of the matrices on the left hand side reveals that they and therefor also their product are all upper triangular. Similarly, the left hand side is a lower triangular matrix and the equality of the two implies that they must both be diagonal. Cancelling D and rearranging the equation yields $\hat{U} = (\hat{L}^{-1} \hat{U}^t) \hat{L}^t$ where $\hat{L}^{-1} \hat{U}^t$ is diagonal. This shows that the rows of $\hat{U}$ are just multiples of the rows of $\hat{L}^t$. However, we know that the diagonal entries of $\hat{U}$ and $\hat{L}$ are the same. Thus, $\hat{L}^{-1} \hat{U}^t$ is the identity and $\hat{L} = \hat{U}^t$.

We now write $R = \hat{L}^t = \hat{U}$. The proof of [15, Theorem 8] shows $A = \Theta D^{-1} R$ and $\Theta^t \Theta = \hat{U} (D \hat{L}^{-1})^t$. Expanding the last expression and using the definition of R gives then $\Theta^t \Theta = R R^{-1} D = D$. □

We now give an explicit expression of the last column of Θ, showing the common factor of $\det A$.

THEOREM 2. *With $A \in \mathbb{D}^{n \times n}$ and Θ as in theorem 1, we have for all $i = 1, \ldots, n$ that*

$$\Theta_{in} = (-1)^{n+i} \det A_{in} \det A$$

where $\det A_{in}$ *is the* (i,n) *minor of* A.

PROOF. We use the notation from the proof of theorem 1. From $\Theta D^{-1}\hat{L}^t = A$ we obtain

$$\Theta^t A = D\hat{L}^{-1}A^t A = D\hat{L}^{-1}(\hat{L}D^{-1}\hat{U}) = \hat{U}.$$

Thus, since A has full rank, $\Theta^t = \hat{U}A^{-1}$ or, equivalently,

$$\Theta = (\hat{U}A^{-1})^t = (A^{-1})^t\hat{U}^t = (\det A)^{-1}(\operatorname{adj} A)^t\hat{U}^t$$

where $\operatorname{adj} A$ is the adjugate matrix of A. Since $\hat{U}^t$ is a lower triangular matrix with $\det A^t A = (\det A)^2$ at position (n,n), the claim follows. □

THEOREM 3. *Given a square matrix* A*, a reduced fraction-free* QR *decomposition is given by* $A = \hat{\Theta}\hat{D}^{-1}\hat{R}$*, where* $S = \operatorname{diag}(1, 1, \ldots, \det A)$ *and* $\hat{\Theta} = \Theta S^{-1}$*, and* $\hat{R} = S^{-1}R$*. In addition,* $\hat{D} = S^{-1}DS^{-1} = \hat{\Theta}^t\hat{\Theta}$.

PROOF. By theorem 2, ΘS^{-1} is an exact division. The theorem follows from $A = \Theta S^{-1}SD^{-1}SS^{-1}R$. □

As an example we consider the 4-by-4 integer matrix

$$A = \begin{pmatrix} -62 & 21 & 64 & -96 \\ 38 & 18 & 31 & 56 \\ -59 & -86 & 19 & 2 \\ 40 & -91 & -62 & 9 \end{pmatrix}.$$

Computing the QR decomposition with theorem 1 yields

$$\Theta = \begin{pmatrix} -62 & 268341 & 2658137038 & -23374523883001 \\ 38 & 155634 & 8243861790 & 3112061098992 \\ -59 & -843590 & 2460946816 & 14218033256642 \\ 40 & -976219 & -81659738 & -18215371009147 \end{pmatrix}$$

$D = \operatorname{diag}(10369, 1760876458298, 81089877269400184044,$
$109000550172869435495496 5838)$,

$R =$

$$\begin{pmatrix} 10369 & 816 & -6391 & 8322 \\ 0 & 169821242 & 66495846 & -27518383 \\ 0 & 0 & 477501379182 & 210662060582 \\ 0 & 0 & 0 & 2282727441742609 \end{pmatrix}.$$

We can now check that indeed $\det A = 47777897$ divides the last column of Θ.

Cancelling $\det A$ from the last column of Θ and the last entry of R as well as reducing D accordingly leads to the much simpler output

$$\Theta = \begin{pmatrix} -62 & 268341 & 2658137038 & -489233 \\ 38 & 155634 & 8243861790 & 65136 \\ -59 & -843590 & 2460946816 & 297586 \\ 40 & -976219 & -81659738 & -381251 \end{pmatrix},$$

$D = \operatorname{diag}(10369, 1760876458298, 81089877269400184044,$
$477501379182)$ and

$$R = \begin{pmatrix} 10369 & 816 & -6391 & 8322 \\ 0 & 169821242 & 66495846 & -27518383 \\ 0 & 0 & 477501379182 & 210662060582 \\ 0 & 0 & 0 & 47777897 \end{pmatrix}$$

4. COMMON FACTORS IN LU

Given a matrix A over an integral domain $\mathbb{D}$, we consider the fraction-free decomposition $A = LD^{-1}U$. It is clear that if the elements in a column of L or a row of U possess a common GCD, then that factor can be removed, reducing the size of the matrix elements. We identify 3 sources of common GCDs.

4.1 Input data

The initial matrix may contain one or more rows having a common GCD, usually because of modelling choices made by the user. Standard Gaussian elimination will then transfer the common factor into all subsequent rows. If several rows have different GCDs, then all GCDs accumulate in subsequent rows.

4.2 LU and the Smith Form

The following theorem links the Smith normal form of a given matrix with factors appearing in the LU decomposition.

THEOREM 4. *Let* $A \in \mathbb{D}^{n\times n}$ *have the Smith normal form* $S = \operatorname{diag}(d_1, \ldots, d_n)$ *where* $d_1, \ldots, d_n \in \mathbb{D}$*. Moreover, let* $A = LD^{-1}U$ *be an* $LD^{-1}U$ *decomposition of* A*. Then for* $k = 1, \ldots, n$

$$d_k^* = \prod_{j=1}^{k} d_j \mid U_{k,*} \quad \textit{and} \quad d_k^* \mid L_{*,k}.$$

REMARK 1. *The values* $d_1^*, \ldots, d_n^*$ *are known as the* determinantal divisors *of* A.

PROOF. According to [11, II.15], the diagonal entries of the Smith form are quotients of the determinantal divisors, i. e., $d_1^* = d_1$ and $d_k = d_k^*/d_{k-1}^*$ for $k = 2, \ldots, n$. Moreover, d_k^* is the greatest common divisor of all k-by-k minors of A for each $k = 1, \ldots, n$. Thus, we only have to prove that the entries of the k^{th} row of U are k-by-k minors of A. However, this follows from [4, Eqns (9.8), (9.12)], since the k^{th} row of U are just

$$\det \begin{pmatrix} A_{11} & \cdots & A_{1k} & A_{1j} \\ \vdots & & \vdots & \vdots \\ A_{k1} & \cdots & A_{kk} & A_{kj} \end{pmatrix} \quad \text{where} \quad j = 1, \ldots, k.$$

Similarly, following the algorithm in [7], we see that the columns of L are just made up by copying entries from the columns of U during the reduction. More precisely, the k^{th} column of L will have the entries $a_{1k}^{(k-1)}, \ldots, a_{nk}^{(k-1)}$ (using the notation of [4]). But these are again just k-by-k minors of A. □

We give an example using the domain $\mathbb{Q}[x]$. Let A be the polynomial matrix

$$\begin{pmatrix} -\frac{3}{2} & -x^3 + 5x^2 + 3x - \frac{9}{2} & x^2 + x & \frac{1}{2}x^3 - x^2 \\ -3 & -2x^3 + 10x^2 + 5x - 9 & 2x^2 + 2x & x^3 - 2x^2 \\ \frac{1}{2} & x^3 + \frac{3}{2} & 0 & -\frac{1}{2}x^3 \\ -\frac{1}{2} & -x - \frac{3}{2} & 0 & \frac{1}{2}x \end{pmatrix}.$$

The Smith normal form S of A is

$$\operatorname{diag}(1, x, x(x+1), x(x+1)(x-1))$$

and thus its determinantal divisors are $d_1^* = 1$, $d_2^* = x$, $d_3^* = x^2(x+1)$ and $d_4^* = x^3(x+1)^2(x-1)$. Computing the $LD^{-1}U$ decomposition of A yields $A = LD^{-1}U$ where L is

$$\begin{pmatrix} -\frac{3}{2} & 0 & 0 & 0 \\ -3 & \frac{3}{2}x & 0 & 0 \\ \frac{1}{2} & -x^3 - \frac{5}{2}x^2 - \frac{3}{2}x & \frac{1}{2}x^3 + \frac{1}{2}x^2 & 0 \\ -\frac{1}{2} & -\frac{1}{2}x^3 + \frac{5}{2}x^2 + 3x & -\frac{1}{2}x^3 - \frac{1}{2}x^2 & -\frac{1}{4}x^6 - \frac{1}{4}x^5 + \frac{1}{4}x^4 + \frac{1}{4}x^3 \end{pmatrix},$$

$D = \operatorname{diag}(-3/2, -9/4x, 3/4x^4 + 3/4x^3, -1/8x^9 - 1/4x^8 + 1/4x^6 + 1/8x^5)$, U is

$$\begin{pmatrix} -\frac{3}{2} & -x^3 + 5x^2 + 3x - \frac{9}{2} & x^2 + x & \frac{1}{2}x^3 - x^2 \\ 0 & \frac{3}{2}x & 0 & 0 \\ 0 & 0 & \frac{1}{2}x^3 + \frac{1}{2}x^2 & -\frac{1}{2}x^4 - \frac{1}{2}x^3 \\ 0 & 0 & 0 & -\frac{1}{4}x^6 - \frac{1}{4}x^5 + \frac{1}{4}x^4 + \frac{1}{4}x^3 \end{pmatrix}.$$

Computing the column factors of L and the row factors of U yields 1, x, $x^2(x+1)$ and $x^3(x-1)(x+1)^2$, i. e., exactly the determinantal divisors. In general, there could be other factors as well.

4.3 Statistical effects

Suppose that during Bareiss's algorithm after $k-1$ iterations we have reached the following state

$$A^{(k-1)} = \begin{pmatrix} U & \underline{*} & \underline{*} & * \\ \overline{0} & p & * & \overline{*} \\ 0 & 0 & a & \overline{v} \\ 0 & 0 & b & \overline{w} \\ \mathbf{0} & \underline{0} & \underline{*} & * \end{pmatrix},$$

where U is an upper triangular matrix, $p, a, b \in \mathbb{D}$, $\overline{v}, \overline{w} \in \mathbb{D}^{1\times n-k-1}$ and the other overlined quantities are row vectors and the underlined quantities are column vectors. Assume that $a \neq 0$ and that we choose it as a pivot. Continuing the computations we now eliminate b (and the entries below) by cross-multiplication

$$A^{(k-1)} \rightsquigarrow \begin{pmatrix} U & \underline{*} & \underline{*} & * \\ \overline{0} & p & * & \overline{*} \\ 0 & 0 & a & \overline{v} \\ 0 & 0 & 0 & a\overline{w} - b\overline{v} \\ \mathbf{0} & \underline{0} & \underline{0} & * \end{pmatrix}.$$

Here, we can see that any common factor of a and b will be a factor of every entry in that row, i. e., $\gcd(a,b) \mid a\overline{w} - b\overline{v}$. However, we still have to carry out the exact division step. This leads to

$$A^{(k-1)} \rightsquigarrow \begin{pmatrix} U & \underline{*} & \underline{*} & * \\ \overline{0} & p & * & \overline{*} \\ 0 & 0 & a & \overline{v} \\ 0 & 0 & 0 & \frac{1}{p}(a\overline{w} - b\overline{v}) \\ \mathbf{0} & \underline{0} & \underline{0} & * \end{pmatrix} = A^{(k)}.$$

The division by p is exact. Some of the factors in p might be factors of a or b while others are hidden in $\overline{v}$ or $\overline{w}$. However, every common factor of a and b which is not also a factor of p will still be a common factor of the resulting row. In other words,

$$\frac{\gcd(a,b)}{\gcd(a,b,p)} \;\Big|\; \frac{1}{p}(a\overline{w} - b\overline{v}).$$

In fact, the factors do not need to be tracked during the $LD^{-1}U$ reduction but can be computed afterwards: All the necessary entries a, b and p of $A^{(k-1)}$ will end up as entries of L. More precisely, we will have $p = L_{k-1,k-1}$, $a = L_{k,k}$ and $b = L_{k+1,k}$.

If $\mathbb{D}$ are the integers, then the probability that the quotient $\gcd(a,b)/\gcd(p,a,b) \neq 1$, i.e. nontrivial, for random a, b, p equals $1 - 6\zeta(3)/\pi^2 \approx 26.92\%$ [6, 14]. Thus, for integer matrices these factors occur with a high enough frequency to suggest we care about them. In our experiments we saw that independently of the size of the input matrix this method could detect about 40.17% of all the common prime row factors occurring in U.[1]

As an example consider the matrix

$$A = \begin{pmatrix} 0 & -18 & -92 & -25 & -60 \\ 49 & -77 & 66 & 45 & 8 \\ 18 & 31 & 69 & -81 & 51 \\ -58 & 41 & 22 & 37 & -97 \\ -77 & -52 & 48 & -19 & -10 \end{pmatrix}.$$

This matrix has a $LD^{-1}U$ decomposition with

$$L = \begin{pmatrix} 8 & 0 & 0 & 0 & 0 \\ -10 & -126 & 0 & 0 & 0 \\ 51 & -2355 & 134076 & 0 & 0 \\ -97 & 4289 & -233176 & -28490930 & 0 \\ -60 & 2940 & -148890 & -53377713 & 11988124645 \end{pmatrix}$$

and

$$U = \begin{pmatrix} 8 & 49 & 45 & -77 & 66 \\ 0 & -126 & 298 & -1186 & 1044 \\ 0 & 0 & 134076 & -414885 & 351648 \\ 0 & 0 & 0 & -28490930 & 55072620 \\ 0 & 0 & 0 & 0 & 11988124645 \end{pmatrix}.$$

The method outlined above correctly predicts the common factor 2 in the second row, the factor 3 in the third row and the factor 2 in the fourth row. However, it does not detect the additional factor 5 in the fourth row.

There is another way in which common factors in integer matrices can arise: Let d be any number. Then for random a, b the probability that $d \mid a + b$ is $1/d$. That means that if $v, w \in \mathbb{Z}^{1\times n}$ are vectors, then $d \mid v + w$ with a probability of $1/d^n$. This effect is noticable in particular for small numbers like $d = 2, 3$ and in the last iterations of the $LD^{-1}U$ decomposition when the number of non-zero entries in the rows has shrunk. For instance, in the second last iterations we only have three rows with at most three non-zero entries each. Moreover, we know that the first non-zero entries of the rows cancel during cross-multiplication. Thus, a factor of 2 appears with a probability of 25% in one of those rows, a factor of 3 with a probability of 11.11%. In the example above, the probability for the factor 5 to appear in the fourth row was 4%.

In a manner similar to theorem 3, we can cancel all factors which we find from the final output:

THEOREM 5. *Given a matrix $A \in \mathbb{D}^{m\times n}$ with rank r and its decomposition $A = P_w L D^{-1} U P_c$, if $D_U = \operatorname{diag}(d_1, \ldots, d_r)$ is a diagonal matrix with $d_k \mid \gcd(U_{k,*})$, then setting $\hat{U} = D_U^{-1} U$ and $\hat{D} = D D_U^{-1}$ where both matrices are fraction-free we have the decomposition $A = P_w L \hat{D}^{-1} \hat{U} P_c$.*

PROOF. By [7, Theorem 2] the diagonal entries of U are the pivots chosen during the decomposition and they also divide the diagonal entries of D. Thus, any common divisor of $U_{k,*}$ will also divide D_{kk} and therefor both $\hat{U}$ and $\hat{D}$ are fraction-free. We can easily check that $A = P_w L D^{-1} D_U D_U^{-1} U = P_w L \hat{D}^{-1} \hat{U} P_c$. □

[1] This experiment was carried out with random square matrices A of sizes between 5-by-5 and 125-by-125. We decomposed A into $P_w L D^{-1} U P_c$ and then computed the number of predicted prime factors in U and related that to the number of actual prime factors. We did not consider the last row of U since this contains only the determinant.

REMARK 2. *If we can find common column factors of L we can cancel them in the same way. However, if we have already cancelled factors from U, then there is no guarantee that $d \mid L_{*,k}$ implies $d \mid \hat{D}_{kk}$. Thus, in general we can only cancel* $\gcd(d, \hat{D}_{kk})$ *from* $L_{*,k}$.

5. PIVOTING STRATEGIES FOR LU

Our pivoting strategies are all based on full pivoting, which is already implied by the definition of the form. We define a number of pivoting strategies.

Largest We select the largest pivot according to an appropriate metric. Metrics were the absolute value for integer matrices and the degree as well as the height for matrices univariate polynomials.

Smallest Here we select the smallest pivot according to the same metrics as above.

First We select the first non-zero pivot.

Factors With this strategy we select the pivot which has the least number of prime factors counted with multiplicity.

Of course, the "factors" strategy is not viable in practice since the factorisation is much too costly. However, it does provide interesting theoretical insight.

In contrast to floating point calculations, accuracy of the result is not an issue, and we consider instead the size of the elements in the matrices generated, and any impact on the efficiency of the computation. By size we examine the following

Digits For integer matrices or matrices we count the total number of base-10 digits needed to represent it. We also use this measurement for matrices with rational number entries where we simply add up the digits of the numerators and the denominators.

Terms For univariate polynomial matrices we count the total number of non-zero terms in the fully expanded representation of the entries.

Height As another metric for polynomial matrices we use the maximal height of its entries.

Factors For both integer and polynomial matrices we measure the total number of row factors. Here, we compute the greatest common divisor of each row and count the number of prime factors with multiplicity. The number of factors for ech row is then added up.

Note that the measured quantities do solely depend on the output. In particular do they not depend on how the programme handles its memory during the computations. Also note that the measurements are chosen in such a way that they are independent of the internal representation of the data. For instance, every programme has to store all the digits of the output matrices somehow.

The experiments included in this paper were all carried out with MAPLE. We use our own implementation of the $LD^{-1}U$ decomposition which closely follows [7]. For each experiment we generated random matrices A of different sizes and then performed the decomposition $A = P_w L D^{-1} U P_c$ using the strategies described above. That is, each random matrix A was decomposed with each of the strategies. We then applied the applicable measurements. In the end we computed the mean value of all the results. More precise description of the experiments follow below.

n	digits			row factors		
	smallest	largest	factors	smallest	largest	factors
5	78.13	101.74	85.13	7.58	8.01	5.74
10	503.72	678.40	569.40	11.65	12.80	6.44
15	1625.08	2130.83	1833.94	17.17	17.95	7.77
20	3832.33	4888.83	4297.05	21.38	22.88	7.98
25	7533.28	9365.39	8316.27	26.06	27.92	8.26

Table 2: Output sizes for different pivoting strategies for integer matrices. The table compares the average number of digits and the number of row common factors of U for random n-by-n integer matrices A as input using the "smallest", "largest" and "factors" pivoting strategies.

n	smallest degree	largest degree	height
5	83.07	106.80	91.91
10	532.45	698.15	609.13
15	1696.09	2154.53	1946.16
20	3932.09	4860.95	4504.71

Table 3: Output sizes for different pivoting strategies for polynomial matrices. The table compares the number of terms of U for random n-by-n input matrices A using the "smallest degree", "largest degree" and "smallest height" pivoting strategies.

For table 2 we generated three hundred integer matrices for each size. The entries where between -11^3 and 11^3. Also in order to be closer to real world problems, we made sure that the sizes of the entries in our matrices varied widely with less than 25% of the entries reaching maximal size. Table 2 shows the number of digits and the number of row factors of U where the decompositions are done using the "smallest", "largest" and "factors" strategies described at the beginning of this section.

Table 3 shows a similar experiment for matrices of univariate polynomials. We compare the strategies of choosing the pivot with the smallest degree versus choosing the largest degree and choosing the smallest height. The matrices A contained random polynomials with integer coefficients between -100 and 100 and degree at most 3. During the same experiment we also measured the number of row factors and the height of U but we did not find a significant difference between the different strategies.

6. SOLVING

In this section we detail a method for solving linear systems in such a way that fractions are delayed until the final output.

Let $A \in \mathbb{D}^{m \times n}$ and $b \in \mathbb{D}^m$. We wish to solve the system $Ax = b$, seeking solutions x with entries in the field of fractions of $\mathbb{D}$. First, apply the $LD^{-1}U$ decomposition as in [7]. We obtain

$$DL^{-1}P_w^t A = \begin{pmatrix} V \\ W \end{pmatrix} A = \begin{pmatrix} U & B \\ 0 & 0 \end{pmatrix} P_c \quad \text{and} \quad P_c x = \begin{pmatrix} y \\ z \end{pmatrix},$$

where all (sub) matrices have entries in $\mathbb{D}$, U is an r-by-r, regular and upper triangular matrix, r is the rank of A and where y has dimension r. Then $Ax = b$ if and only if $Wb = 0$

and $Uy + Bz = Vb$.

Now, perform a second $LD^{-1}U$ decomposition on U (pivoting is not needed as all diagonal entries of U are non-zero), working from the bottom to the top, and from right to left[2]. This will compute a regular $X \in \mathbb{D}^{r\times r}$ such that $XU = \Delta$ is a diagonal matrix. Then $Ax = b$ if and only if $Wb = 0$ and $\Delta y + XBz = XVb$.

Assume now that the compatibility condition $Wb = 0$ is fulfilled. In order to compute a particular solution x_0 of the system $Ax = b$, we can simply choose

$$x_0 = P_c^{-1}\begin{pmatrix}\Delta^{-1}XVb\\ 0\end{pmatrix} = \tilde{\Delta}^{-1}Sb \quad \text{where} \quad S = P_c\begin{pmatrix}XV\\ 0\end{pmatrix}$$

and where $\tilde{\Delta} = P_c\,\mathrm{diag}(\Delta, \mathbf{1})P_c$ is a diagonal matrix with entries in $\mathbb{D}$.

Moreover, we can compute the nullspace of A in the following way: If

$$x \in \mathrm{colspace}\, P_c^{-1}\begin{pmatrix}-\Delta^{-1}\,XB\\ \mathbf{1}_{n-r}\end{pmatrix},$$

then we can easily check that $Ax = 0$. Since the $n - r$ columns of the matrix spanning the space are clearly linearly independent, it follows that this is already the entire nullspace of A. Thus, setting

$$K = P_c\begin{pmatrix}-XB\\ \mathbf{1}\end{pmatrix},$$

we see the nullspace of A is colspace $\tilde{\Delta}^{-1}K$, with $\tilde{\Delta}$ as defined above.

Note that S and K are both matrices over $\mathbb{D}$. Thus, the particular solution and the nullspace are both computed in a fraction-free way. Moreover, neither of the matrices depends on the right hand side b. Consequently, after computing W, S, $\tilde{\Delta}$ and K, we can solve the system $Ax = b$ for arbitrary b by just checking whether $Wb = 0$ and then computing $x_0 = \tilde{\Delta}^{-1}Sb$.

We summarise the method as follows:

ALGORITHM 1. **Input:** *A matrix $A \in \mathbb{D}^{m\times n}$.*

Output: *Matrices W, S, and K with entries in $\mathbb{D}$ and a diagonal matrix $\tilde{\Delta}$ with entries in $\mathbb{D}$ such that for any $b \in \mathbb{D}^m$ if the compatibility condition $Wb = 0$ is met, then the system $Ax = b$ has the solution set $\tilde{\Delta}^{-1}Sb +$ colspace $\tilde{\Delta}^{-1}K$.*

Steps:

1. *Apply the $LD^{-1}U$ decomposition to obtain*
$$DL^{-1}P_w^tA = \begin{pmatrix}V\\ W\end{pmatrix}A = \begin{pmatrix}U & B\\ 0 & 0\end{pmatrix}$$
where U is upper triangular.
2. *Use a backwards $LD^{-1}U$ decomposition on U to obtain a matrix X such that diagonal $XU = \Delta$ is a diagonal matrix.*
3. *Let*
$$S = P_c\begin{pmatrix}XV\\ 0\end{pmatrix}, \quad K = P_c\begin{pmatrix}-XB\\ \mathbf{1}\end{pmatrix}$$
and $\tilde{\Delta} = P_c\,\mathrm{diag}(\Delta, \mathbf{1})P_c$.

[2]More formally, let Π be the matrix of the permutation which maps i to $r+1-i$ and decompose $\Pi U\Pi$ in the normal way applying the same permutations to the result.

As an example we consider the matrix

$$A = \begin{pmatrix}-370 & -62 & -101 & -3\\ -708 & -120 & -193 & -5\\ -304 & -50 & -83 & -3\\ -1962 & -336 & -534 & -12\end{pmatrix}$$

and examine the two systems below for solutions.

$$Ax = \begin{pmatrix}1\\0\\0\\1\end{pmatrix} = b_1 \quad \text{and} \quad Ax = \begin{pmatrix}0\\0\\1\\1\end{pmatrix} = b_2.$$

Following algorithm 1, we first compute

$$\left(\begin{array}{cccc}1 & 0 & 0 & 0\\ 3 & 0 & -3 & 0\\ 110 & -36 & -50 & 0\\ \hline 7 & -6 & -1 & 1\end{array}\right)A = \left(\begin{array}{ccc|c}-3 & -62 & -101 & -370\\ 0 & -36 & -54 & -198\\ 0 & 0 & -12 & -12\\ \hline 0 & 0 & 0 & 0\end{array}\right)P_c$$

where P_c represents the permutation $(1\quad 4)$; and use this to define the matrices V, W, U and B. Next, we compute

$$X = \begin{pmatrix}432 & -744 & -288\\ 0 & -12 & 54\\ 0 & 0 & 1\end{pmatrix}$$

and $XU = \mathrm{diag}(-1296, 432, -12) = \Delta$. This leads to

$$S = \begin{pmatrix}0 & 0 & 0 & 0\\ 5904 & -1944 & -2664 & 0\\ 110 & -36 & -50 & 0\\ -33480 & 10368 & 16632 & 0\end{pmatrix}, \quad K = \begin{pmatrix}1\\ -1728\\ 12\\ 9072\end{pmatrix}$$

and $\tilde{\Delta} = \mathrm{diag}(1, 432, -12, -1296)$.

We can check that $Wb_1 = 8 \neq 0$. Consequently, the system $Ax = b_1$ does not have a solution. On the other hand, $Wb_2 = 0$ and the solution set for $Ax = b_2$ is

$$\tilde{\Delta}^{-1}Sb + \mathrm{colspace}\,\tilde{\Delta}^{-1}K = \begin{pmatrix}0\\ -37/6\\ 25/6\\ -77/6\end{pmatrix} + \mathrm{colspace}\begin{pmatrix}1\\ -4\\ -1\\ -7\end{pmatrix}.$$

7. CONCLUSIONS

We have shown that fraction-free LU and QR decompositions can contain significant common factors, and we have shown how these can be beneficially removed to obtain more compact decompositions. Moreover, their removal makes the decomposition unique.

We considered removing the common factors as soon as they can be detected during the computation of the decompositions. This would require either discovering the GCDs by direct computation, or by predicting them by different, preferably simpler, computations. Although we have displayed here mechanisms that generate common factors, and which lend themselves to predictions through relatively simple calculations, there are other mechanisms which we have not discussed. These require more extensive computations to predict, and quickly leave the realm of reasonable strategies. Therefore we have concluded that it is most sensible to leave common factor identification to the final stage of decomposition.

We hope that reduced decompositions can be implemented as the standard form in future computer-algebra systems.

8. ACKNOWLEDGMENTS

This work was supported in part by the Austrian Science Fund (FWF) grant SFB50 (F5009-N15).

We would like to thank Prof. Kevin G. Hare and Univ.-Doz. Dr. Arne Winterhof for helpful.

9. REFERENCES

[1] E. H. Bareiss. Sylvester's identity and multistep integer-preserving Gaussian elimination. *Mathematics of Computation*, 22(103):565 – 578, 1968.

[2] J.-G. Dumas, C. Pernet, and Z. Sultan. Computing the rank profile matrix. In D. Robertz, editor, *Proceedings of the 2015 International Symposium on Symbolic and Algebraic Computation, ISSAC'15*, pages 149–156. ACM, ACM Press, 2015.

[3] Ú. Erlingsson, E. Kaltofen, and D. Musser. Generic Gram—Schmidt orthogonalization by exact division. In *International Symposium on Symbolic and Algebraic Computation*, pages 275–282. ACM press, 1996.

[4] K. Geddes, G. Labahn, and S. Czapor. *Algorithms for Computer Algebra*. Kluwer, 1992.

[5] M. W. Giesbrecht and A. Storjohann. Computing rational forms of integer matrices. *Journal of Symbolic Computation*, 34(3):157–172, 2002.

[6] K. G. Hare. Personal Communication.

[7] D. J. Jeffrey. LU factoring of non-invertible matrices. *Comm. Comp. Alg.*, 44(171):1–8, 2010.

[8] E. Kaltofen and G. Yuhasz. A fraction free matrix Berlekamp/Massey algorithm. *Linear Algebra and Applications*, 439(9):2515–2526, 2013.

[9] H. R. Lee and B. D. Saunders. Fraction free Gaussian elimination for sparse matrices. *J. Symbolic Computation*, 19:393–402, 1995.

[10] G. C. Nakos, P. R. Turner, and R. M. Williams. Fraction-free algorithms for linear and polynomial equations. *SIGSAM Bull.*, 31(3):11–19, 1997.

[11] M. Newman. *Integral Matrices*, volume 45 of *Pure and Applied Mathematics*. Academic Press, New York, 1972.

[12] C. Pauderis and A. Storjohann. Computing the invariant structure of integer matrices: fast algorithms into practice. In M. Kauers, editor, *Proceedings of the International Symposium on Symbolic and Algebraic Computation, ISSAC'13*. ACM Press, 2013.

[13] L. Pursell and S. Y. Trimble. Gram-Schmidt orthogonalization by Gaussian elimination. *American Math. Monthly*, 98(6):544–549, 1991.

[14] A. Winterhof. Personal Communication.

[15] W. Zhou and D. J. Jeffrey. Fraction-free matrix factors: new forms for LU and QR factors. *Frontiers of Computer Science in China*, 2(1):67–80, 2008.